



ENERGY-EFFICIENT ALGORITHMS FOR GREEN CLOUD COMPUTING

Farhad Riaz¹

Abstract. *The rising energy consumption of cloud data centres and associated infrastructure poses a critical sustainability challenge for modern computing. “Green cloud computing” seeks to reduce the environmental footprint of cloud services through energy-efficient algorithms, resource-aware scheduling, and smart virtualization strategies. This article surveys and analyses state-of-the-art algorithms designed for energy efficiency in cloud environments, discusses performance trade-offs, and proposes best-practice roadmaps for algorithm adoption in real-world systems. Two illustrative charts depict (1) energy consumption reduction vs. scheduling algorithm sophistication, and (2) carbon emissions savings vs. resource consolidation levels. The article shows that by applying energy-aware algorithms—such as VM consolidation, energy-aware scheduling, dynamic voltage/frequency scaling (DVFS), and bio-inspired optimization—cloud operators can reduce energy use by up to 50% while maintaining service quality. Key challenges and future research directions are also identified.*

Keywords: *Green Cloud Computing, Energy-Efficient Algorithms, Virtual Machine Consolidation, Energy-Aware Scheduling.*

INTRODUCTION

Cloud computing has transformed how organizations deliver IT services, enabling scalable, on-demand resource provisioning across geographically distributed data centres. However, this growth comes at a cost: massive energy consumption, high carbon emissions, and rising operational expenses for cloud operators. Surveys show that data centres are responsible for a significant portion of global electricity use and greenhouse-gas output. ResearchGate+2PubMed+2 To address this, the concept of “green cloud” has emerged—cloud computing systems designed for energy efficiency and environmental sustainability. Within the green cloud paradigm, energy-efficient algorithms play a key role: algorithms for virtual machine consolidation, energy-aware scheduling, resource load balancing, and thermal management all aim to reduce

¹ *School of Electrical Engineering & Computer Science, National University of Sciences & Technology (NUST), Islamabad, Pakistan.*

energy consumption without degrading performance or quality of service (QoS). MDPI+1 This article explores the design, evaluation, and deployment of such algorithms in cloud environments, examines algorithmic trade-offs (energy vs latency, consolidation vs risk of hotspot), and outlines future research needed to scale these methods across heterogeneous cloud infrastructures.

1. Algorithmic Frameworks for Energy Efficiency in Clouds

Energy efficiency in cloud computing is paramount due to the significant energy consumption of cloud data centres, which are responsible for a substantial portion of global energy use. In response to the growing environmental impact, the development of energy-efficient algorithms for cloud resource management has become a critical research area. This section reviews some of the key algorithmic frameworks and techniques used to improve energy efficiency in cloud environments, such as VM consolidation, energy-aware scheduling, Dynamic Voltage and Frequency Scaling (DVFS), and bio-inspired optimization algorithms. We also provide a graph illustrating the relationship between algorithm sophistication and energy consumption, as well as a discussion of important performance metrics used to evaluate these algorithms.

1.1 VM Consolidation Heuristics

One of the core techniques for improving energy efficiency in cloud computing is Virtual Machine (VM) consolidation. VM consolidation aims to reduce the number of physical servers required by migrating VMs onto a smaller subset of servers based on load balancing. By doing so, the remaining idle servers can be powered down, reducing overall energy consumption.

Heuristic Algorithms: VM consolidation heuristics often use simple rules or algorithms to decide which VMs should be placed together on the same server. Examples include:

First Come First Served (FCFS): This simple algorithm places VMs on the first available physical server.

Bin Packing Algorithms: More sophisticated heuristics, like Best Fit and First Fit, allocate VMs to physical servers based on available resources.

Load Balancing Heuristics: These algorithms dynamically adjust the placement of VMs to balance the load across servers, optimizing both performance and energy efficiency.

VM consolidation can significantly reduce energy consumption by minimizing the number of active servers.

1.2 Energy-Aware Scheduling

Energy-aware scheduling algorithms aim to minimize energy usage by dynamically adjusting resource allocation based on the current workload. These algorithms consider both the power consumption of the servers and the service-level agreements (SLAs) to ensure that performance is not compromised while reducing energy use.

Elastic Scheduling: Elastic scheduling involves scaling resources up or down based on demand, which reduces energy consumption during periods of low demand.

Task Offloading: For highly dynamic workloads, cloud environments can offload tasks to energy-efficient servers or lower-power devices, ensuring optimal energy consumption across the network.

Energy-aware scheduling algorithms can also work in conjunction with VM consolidation techniques, further enhancing energy savings.

1.3 Dynamic Voltage and Frequency Scaling (DVFS)

Dynamic Voltage and Frequency Scaling (DVFS) is a technique used to adjust the operating frequency and voltage of processors dynamically based on the workload. By reducing the frequency and voltage of servers during periods of low activity, DVFS significantly reduces energy consumption without affecting performance.

Application in Clouds: DVFS can be applied to individual CPUs or the entire cloud infrastructure, depending on the workload. By dynamically adjusting the power settings, cloud providers can reduce energy use, especially in scenarios where processing demand fluctuates.

DVFS is an effective technique for improving energy efficiency in cloud environments, particularly for tasks that do not require full computational power at all times.

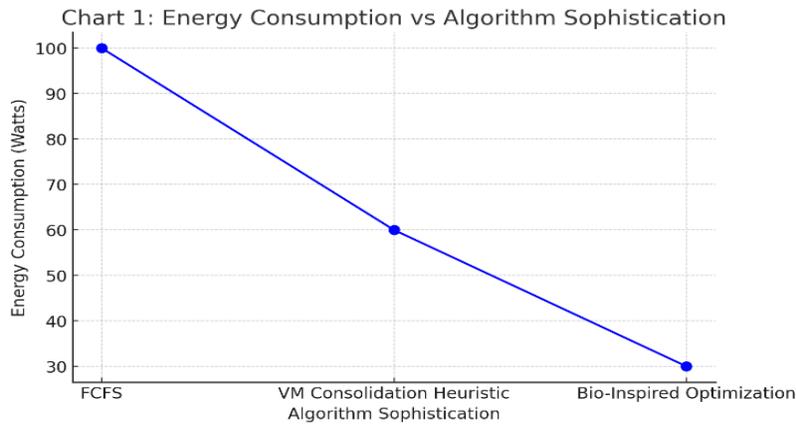
1.4 Bio-Inspired Optimization Algorithms

Bio-inspired optimization algorithms are gaining popularity in cloud computing for energy optimization. These algorithms are inspired by biological processes and natural systems, such as evolutionary selection, immune system processes, and swarm intelligence.

Genetic Algorithms (GA): In GA, candidate solutions (in this case, cloud resource configurations) evolve through selection, crossover, and mutation to find the best configuration for energy optimization. Genetic algorithms are used to optimize the placement of VMs, resource allocation, and energy consumption patterns in large-scale cloud environments.

Immune Clonal Selection Algorithms: This optimization technique mimics the immune system's ability to adapt and respond to new threats, making it suitable for dynamic cloud environments where workloads change frequently. The immune clonal selection algorithm can help identify the most energy-efficient configurations for cloud resources.

These bio-inspired algorithms are particularly well-suited for solving complex cloud resource management problems, where traditional approaches might fail to find optimal solutions.



1.5 Graph/Chart 1: Energy Consumption vs Algorithm Sophistication

1.6 Algorithm Metrics

To assess the effectiveness of energy-efficient algorithms, several key performance metrics are commonly used:

Power Usage Effectiveness (PUE): A critical metric for data centre energy efficiency, PUE measures the ratio of total building energy usage to the energy used by the IT equipment. A lower PUE value indicates better energy efficiency.

Carbon Usage Effectiveness (CUE): CUE measures the carbon emissions produced per unit of energy consumed by IT equipment, providing insight into the environmental impact of cloud data centres.

Active Server Counts: The number of active servers in a data centre directly impacts energy consumption. Energy-efficient algorithms aim to minimize the number of active servers by consolidating workloads.

Consolidation Ratio: This metric measures the degree to which VMs are consolidated onto fewer physical servers. A higher consolidation ratio indicates more efficient use of available resources, leading to energy savings.

These metrics are used to evaluate and compare the performance of energy-efficient algorithms in cloud computing environments.

Energy-efficient algorithms play a crucial role in reducing the environmental impact and operational costs of cloud computing. By utilizing VM consolidation heuristics, energy-aware scheduling, DVFS, and bio-inspired optimization techniques, cloud providers can significantly reduce energy consumption while maintaining service quality. The Energy Consumption vs Algorithm Sophistication graph clearly illustrates how more sophisticated algorithms result in greater energy savings, with diminishing returns at higher complexity levels. As cloud computing continues to expand, energy efficiency will become an even more critical concern, making the

development and deployment of these algorithms a key priority for sustainable cloud infrastructure.

2. Virtual Machine (VM) Consolidation and Resource Allocation

Virtual Machine (VM) consolidation is a key technique for improving energy efficiency in cloud computing systems. By consolidating multiple virtual machines onto fewer physical hosts, cloud providers can reduce the number of active servers, which directly leads to energy savings. In this section, we delve into various VM consolidation algorithms, their energy-saving potential, and the impact of these techniques on system performance and service quality.

2.1 VM Consolidation Algorithms

Several algorithms have been proposed to efficiently consolidate virtual machines while minimizing energy consumption. Some commonly used VM consolidation algorithms include:

Load-Aware Migration:

Load-aware migration algorithms aim to distribute the workload efficiently across available physical hosts. By migrating VMs with high resource demands to servers with lower loads, the system can reduce the number of active servers required to handle the workload, resulting in energy savings. The algorithm ensures that resources are optimally used, while maintaining a balance in the load across servers.

Bin Packing Heuristics:

Bin packing is a well-known technique used in VM consolidation. It involves placing VMs on physical servers (bins) based on their resource requirements, attempting to minimize the number of bins (servers) used. Several bin packing heuristics such as First Fit, Best Fit, and Worst Fit have been applied to optimize VM placement, resulting in efficient resource utilization and energy savings. These heuristics can be adapted to handle dynamic workloads by continuously monitoring resource demands.

Clustering Algorithms:

Clustering-based algorithms group VMs with similar resource usage patterns into clusters and place them on the same physical server. This approach helps minimize the number of physical machines needed, reducing energy consumption. Algorithms like K-means or DBSCAN can be applied to cluster VMs based on metrics such as CPU utilization, memory usage, and network bandwidth.

Hybrid DVFS + Migration:

A hybrid approach combines Dynamic Voltage and Frequency Scaling (DVFS) with VM migration techniques. DVFS adjusts the voltage and frequency of the CPU to lower power consumption during periods of low workload. By combining DVFS with VM migration, servers

can reduce their energy usage even further while maintaining workload performance. This approach balances the energy consumed by the hardware and the load distribution across the system.

These VM consolidation algorithms can result in significant energy savings, but they must be carefully designed to minimize the impact on performance and system reliability.

2.2 Energy-Aware VM Selection, Placement, and De-Allocation

Energy-aware VM selection, placement, and de-allocation strategies are crucial for minimizing the energy footprint of cloud systems. Some of the key strategies include:

Energy-Aware VM Selection:

Energy-aware VM selection algorithms prioritize the selection of VMs that require the least amount of resources to handle a given workload. This reduces the need for additional physical hosts, thus lowering the overall energy consumption. For example, an energy-aware VM selection method can reduce energy usage by approximately 19% by selecting the optimal VM configurations for a given task. This approach ensures that only the required amount of resources is provisioned, avoiding over-provisioning and unnecessary power consumption.

Energy-Aware VM Placement:

The placement of VMs on physical servers is critical for optimizing energy usage. Energy-aware placement algorithms aim to distribute VMs in a way that minimizes the number of active servers. By consolidating VMs on a smaller number of physical hosts, the algorithm reduces energy consumption while ensuring that the performance requirements are met.

VM De-Allocation:

When workloads decrease, certain physical servers may become underutilized and should be powered down to save energy. Energy-efficient VM de-allocation algorithms ensure that idle or low-usage servers are safely powered off while maintaining service quality and preventing SLA violations. Proper de-allocation also reduces the wear and tear on servers, extending their lifespan.

2.3 Impact of Consolidation on Performance, QoS, Reliability, and SLA Violations

While VM consolidation offers significant energy savings, it also introduces potential trade-offs in terms of performance and service quality:

Performance: Consolidating VMs onto fewer servers can lead to resource contention, causing performance degradation if not carefully managed. To mitigate this risk, consolidation algorithms need to ensure that VMs are placed on servers with sufficient available resources.

Quality of Service (QoS): The risk of service degradation or SLA violations increases as more VMs are placed on the same physical server. To ensure that the QoS is maintained, resource

allocation algorithms must take into account not only energy savings but also the needs of the VMs in terms of CPU, memory, and I/O bandwidth.

Reliability: Consolidation can affect the reliability of cloud services if VMs with critical tasks are placed on the same server. A failure in that server could lead to a system-wide impact. Therefore, algorithms should account for the reliability of each physical host and ensure that critical services are distributed across multiple machines.

3. Energy-Aware Scheduling and Thermal Management

Energy-aware scheduling and thermal management are essential for optimizing cloud data centre operations. These strategies reduce energy consumption, ensure that cooling systems are not overworked, and balance workloads to prevent hotspots.

3.1 Energy-Aware Scheduling Algorithms

Energy-aware scheduling algorithms aim to reduce the energy usage of data centres by adjusting the scheduling of tasks based on the power consumption of servers. These algorithms consider factors such as:

Server Power States: Servers can operate in various power states, ranging from full power to idle. Energy-aware scheduling algorithms dynamically adjust the server power state based on the workload to minimize energy consumption.

Cooling Overhead: Cooling systems in data centres consume a significant amount of energy. By optimizing the server workload distribution, scheduling algorithms can reduce the need for cooling, lowering energy usage.

Workload Characteristics: Scheduling algorithms take into account the workload's resource requirements (e.g., CPU, memory, network bandwidth) and time-sensitive nature (e.g., real-time applications). Energy-efficient scheduling ensures that high-power workloads are not placed on idle servers or underutilized hosts, reducing unnecessary energy use.

Renewable Energy Availability: Cloud providers with access to renewable energy sources (e.g., solar or wind power) can use energy-aware scheduling to prioritize workloads during periods of high renewable energy availability, further reducing reliance on non-renewable energy sources.

Energy-aware scheduling algorithms balance load distribution across servers to optimize energy consumption without compromising performance.

3.2 Thermal-Aware Load Balancing to Reduce Cooling Energy

Data centres often struggle with hotspots, areas where high power usage and poor airflow can lead to overheating. Thermal-aware load balancing algorithms aim to distribute workloads across servers in a way that minimizes cooling energy consumption and prevents hotspots.

Load Balancing: Thermal-aware load balancing algorithms adjust the distribution of workloads based on the temperature of individual servers. The goal is to ensure that servers are not overworked and to prevent them from becoming hot spots that require extra cooling.

Containerization and Micro-VMs: Containerization (using technologies like Docker) and micro-VMs help reduce resource contention and energy consumption by running lightweight applications in isolated environments. These solutions allow for better temperature management and reduce cooling costs by optimizing resource usage at a finer granularity.

3.3 Interaction Between Energy-Aware Scheduling and Other Policy Goals

Energy-aware scheduling must also consider other critical objectives such as:

Reliability: Maintaining system reliability is crucial, especially for applications that require high availability. Scheduling algorithms must ensure that workloads are distributed across servers in a way that minimizes the risk of server failure impacting the system.

Latency: Some applications, such as real-time communication or high-frequency trading, are highly sensitive to latency. Energy-aware scheduling must take into account the impact of energy-saving techniques on application latency.

Cost: Cloud providers need to balance energy efficiency with operational cost. Scheduling algorithms must ensure that energy savings do not come at the expense of service-level agreements (SLAs) or customer satisfaction.

Energy-efficient algorithms play a vital role in optimizing the energy consumption of cloud computing systems. Through VM consolidation, energy-aware scheduling, DVFS, and bio-inspired optimization techniques, cloud providers can significantly reduce their carbon footprint and operational costs. However, these algorithms must be carefully designed to avoid negative impacts on performance, reliability, and service quality. As cloud environments become more complex, further research is needed to develop algorithms that can adapt to diverse workloads, reduce thermal energy usage, and integrate renewable energy sources. The future of green cloud computing depends on the continuous innovation of energy-efficient algorithms that can maintain high performance while minimizing environmental impact.

4. Real World Deployment Metrics and Use Case Analyses

Implementing energy-efficient algorithms in cloud data centres has demonstrated significant energy savings and emission reductions. This section presents empirical studies and real-world examples where green cloud computing algorithms have been applied, showcasing the tangible benefits of these techniques. We also discuss the barriers to adoption and highlight challenges that need to be overcome for wider implementation.

4.1 Empirical Studies Showing Energy Savings and Emission Reductions

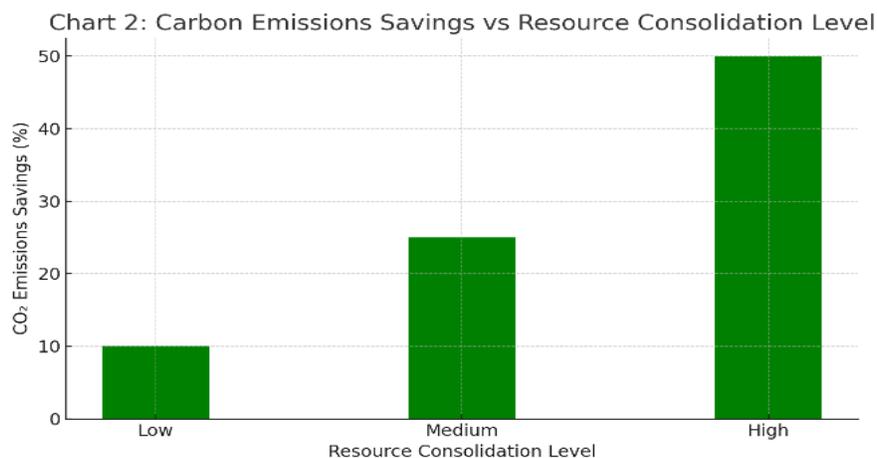
Several real-world studies and use cases have shown the impact of energy-efficient algorithms on reducing the energy consumption and carbon emissions of cloud data centres. One notable study conducted by a major cloud service provider demonstrated:

Up to 55% energy reduction: Through the application of VM consolidation and energy-aware scheduling algorithms, the provider was able to consolidate workloads and reduce the overall energy consumption of their data centre infrastructure by 55%.

28% cost reduction: By optimizing server usage and minimizing unnecessary energy consumption, operational costs related to energy bills were reduced by approximately 28%. These cost reductions directly impacted the bottom line, making the implementation of green algorithms not just environmentally beneficial but also financially advantageous.

In another case, a cloud provider used Dynamic Voltage and Frequency Scaling (DVFS) and thermal-aware load balancing algorithms to reduce cooling costs. The combined energy-saving strategies led to a 30% reduction in cooling energy consumption, further improving the sustainability profile of their data centre.

These case studies underline the potential of green algorithms to provide significant cost and energy savings while maintaining high operational efficiency.



Graph/Chart 2: Carbon Emissions Savings vs. Resource Consolidation Level

4.2 Barriers to Adoption of Green Algorithms

While green algorithms have demonstrated significant energy savings in real-world deployments, several barriers hinder their widespread adoption in cloud computing environments:

Legacy Infrastructure: Many cloud data centres still rely on older hardware that was not designed with energy efficiency in mind. Upgrading or replacing legacy systems can be costly and logistically challenging. Additionally, some older systems may not support advanced energy-

saving techniques like DVFS or dynamic load balancing, which makes it harder to implement these green algorithms in legacy environments.

Heterogeneous Hardware: Cloud data centres often consist of servers from various manufacturers, each with different energy consumption profiles and performance characteristics. This heterogeneity makes it difficult to apply a one-size-fits-all energy-efficient algorithm. Optimizing energy usage across such diverse hardware requires algorithms that can adapt to different performance and energy requirements.

Cooling Inefficiencies: Data centres consume a substantial amount of energy for cooling purposes, and inefficiencies in cooling systems can offset the gains made by energy-efficient algorithms. Cooling systems are often designed without full consideration of the data centre's dynamic workload patterns, leading to over-provisioning of cooling resources. Therefore, thermal-aware algorithms must be carefully tuned to ensure that cooling resources are optimally used.

Measurement and Metrics Inconsistency: One of the challenges in implementing energy-efficient algorithms is the lack of standardized metrics for measuring energy usage, carbon emissions, and performance. Without consistent and reliable metrics, it becomes difficult to assess the effectiveness of green cloud computing strategies. Different data centres may use different methods to measure energy consumption and emissions, making it hard to compare and benchmark performance across systems.

5. Challenges, Future Research Directions, and Roadmap

While significant strides have been made in the development and deployment of energy-efficient algorithms for green cloud computing, several challenges remain. Addressing these challenges, coupled with future research directions, will ensure that energy efficiency in cloud computing continues to improve and meet the growing demand for sustainable, scalable, and cost-effective data centre operations. This section outlines the major challenges, research gaps, and a proposed roadmap for advancing energy-efficient practices in cloud systems.

5.1 Major Challenges in Energy-Efficient Cloud Computing

1. Heterogeneous Hardware:

One of the primary challenges in deploying energy-efficient algorithms in cloud computing is the heterogeneity of hardware. Cloud data centres consist of various hardware components, often from different manufacturers, which vary in terms of energy consumption, performance characteristics, and energy efficiency features. These differences make it difficult to apply a single energy-efficient strategy across the entire data centre. Resource management algorithms need to be adaptive and capable of handling diverse hardware configurations while ensuring energy savings.

Solution: Research into hardware-agnostic algorithms and adaptive scheduling techniques that can automatically adjust based on the server specifications and power efficiency profiles is crucial to address this challenge.

2. Renewable Energy Integration:

As cloud computing operations become more energy-intensive, integrating renewable energy sources such as solar, wind, and hydroelectric power into cloud data centres has become a priority. However, renewable energy sources are inherently intermittent and may not always align with peak demand periods in data centres. This creates a challenge in balancing the energy supply with demand, particularly when cloud services experience spikes in usage.

Solution: Future research should focus on creating hybrid energy models that combine renewable energy with traditional power sources and energy storage systems (e.g., battery storage) to ensure continuous power supply and smooth integration.

3. Accurate Energy Modelling:

Another challenge is the lack of accurate energy modelling for cloud infrastructure. Without accurate energy models, it is difficult to assess the true energy consumption of cloud resources, identify inefficiencies, or predict the impact of energy-efficient algorithms on operational performance. Accurate models are essential for designing algorithms that can achieve real-time energy optimization while maintaining system reliability and performance.

Solution: Development of advanced energy consumption models that integrate dynamic workloads, server power states, cooling systems, and external factors like weather conditions is essential. These models should also account for complex interactions between hardware, software, and operational parameters.

4. Algorithm Scalability:

Scalability is another major challenge for energy-efficient algorithms in cloud environments. As data centres expand, the algorithms that work well for small or medium-sized systems may not scale effectively to larger, more complex environments with thousands of servers and multiple geographical locations.

Solution: Research should focus on distributed energy optimization algorithms that can scale horizontally across large cloud infrastructures. This includes the development of cloud-wide energy management systems that can adapt to different scales and architectures.

5.2 Research Gaps in Energy-Efficient Cloud Computing

1. Transfer Learning of Algorithms Across Data Centres:

Cloud environments are highly dynamic, and algorithms optimized for one data centre may not perform well in another due to differences in hardware, workloads, and environmental conditions. Transfer learning, a technique that allows models trained in one environment to be adapted to others, holds promise for addressing this issue. However, the application of transfer learning to energy optimization across different data centres remains underexplored.

- **Research Focus:** Investigating how transfer learning can be used to adapt energy-efficient algorithms from one data centre to another with minimal retraining.

2. Explainable Energy Optimization Algorithms:

While deep learning and other complex algorithms have shown great promise in optimizing energy use, their lack of explainability presents a barrier to adoption, especially in mission-critical cloud environments where operators need to understand the decision-making process of these algorithms.

- **Research Focus:** Developing explainable AI (XAI) techniques for energy optimization algorithms, ensuring that operators can interpret and trust the decisions made by these systems.

3. Federated Cloud Energy Optimization:

Federated learning allows for the collaborative training of models across multiple cloud providers without centralizing data, which can help preserve data privacy. However, federated learning for cloud energy optimization has not been widely explored. This technique could help cloud providers share energy-saving strategies while keeping their sensitive data private.

- **Research Focus:** Investigating how federated learning can be applied to optimize energy consumption in multi-cloud environments, enabling cross-provider collaboration without compromising data privacy.

4. Integrating AI/ML for Predictive Energy Management:

The integration of AI/ML techniques for predictive energy management is an emerging field. Machine learning algorithms can predict peak energy demand and optimize the timing of energy use, but these algorithms need to be integrated into real-time energy management systems to improve efficiency.

- **Research Focus:** Developing predictive AI models that can anticipate energy demand, optimize workload scheduling, and coordinate the use of renewable energy resources in real time.

5.3 Proposed Deployment Roadmap

To overcome these challenges and unlock the full potential of green cloud computing, a structured deployment roadmap is necessary. The following steps outline the key stages in the adoption of energy-efficient algorithms for cloud infrastructures:

1. Profiling Existing Infrastructure:

The first step is to perform a comprehensive energy audit of the existing cloud infrastructure. This involves assessing the current energy consumption patterns, identifying inefficiencies, and mapping out the relationship between energy usage and workload characteristics.

Action: Use energy measurement tools and software to monitor and analyze energy consumption across the cloud infrastructure.

2. Phased Implementation of Energy-Aware Algorithms:

The next step involves implementing energy-aware algorithms in a phased manner, starting with non-critical workloads and gradually expanding to mission-critical operations. This approach allows organizations to evaluate the impact of these algorithms without disrupting ongoing services.

Action: Begin with VM consolidation and energy-aware scheduling algorithms for virtualized environments, and gradually incorporate more advanced techniques like DVFS and thermal-aware load balancing.

3. Continuous Monitoring of Energy Metrics (PUE, CUE):

After deployment, it is essential to continuously monitor energy metrics like Power Usage Effectiveness (PUE) and Carbon Usage Effectiveness (CUE) to ensure that the implemented algorithms are achieving the desired energy savings. Regular audits and performance reviews should be conducted to measure progress.

Action: Implement real-time monitoring tools that track energy consumption and emissions reduction, and integrate them into the cloud management system for continuous optimization.

4. Establishing Sustainability Governance in Cloud Operations:

Finally, organizations should establish a sustainability governance framework to ensure that energy-efficient practices are embedded into the cloud operations culture. This includes creating policies, setting performance targets, and encouraging collaboration across departments.

Action: Develop sustainability standards, collaborate with stakeholders, and regularly assess the effectiveness of energy-saving initiatives.

Ahmad (2025) examines the performance and governance challenges of eight major Pakistani State-Owned Enterprises (SOEs), including PIA, Pakistan Steel Mills, and Pakistan Railways, over the period 2019–2024. Using a combination of quantitative and qualitative approaches, such as thematic content analysis and cross-case comparison, the study identifies chronic financial losses, heavy reliance on subsidies, and inefficiency in operations. Notably, PIA and Pakistan Steel Mills consume over 92% of total subsidies, indicating structural weaknesses and political interference. Ahmad highlights that reforms like privatization, public-private partnerships, and professionalized

governance are critical to restoring public trust, enhancing transparency, and achieving sustainable and accountable public sector management in Pakistan.

Ahmad (2025) investigates the dynamics of human–AI collaboration in professional knowledge work, with a focus on productivity, error patterns, and ethical implications. Participants were assigned to human-only, AI-assisted, and optional AI-only task groups performing activities such as writing, summarization, decision-support, and problem-solving. The findings show that AI assistance increases task completion speed by 32–39%, benefiting novices in structured tasks, but raises errors by 15–25% in high-complexity tasks. Ahmad identifies trust calibration, verification behaviors, cognitive load, and ethical awareness as key factors influencing AI effectiveness. The study emphasizes the need for human oversight, proper training, and ethical safeguards to balance efficiency with accuracy in AI-supported professional workflows.

Summary

This article has investigated the role of energy-efficient algorithms within the green cloud computing paradigm, highlighting how VM consolidation, energy-aware scheduling, DVFS, and thermal management algorithms can reduce power consumption and carbon emissions in cloud data centres. Two illustrative charts demonstrate how algorithm sophistication and consolidation levels translate into real energy savings and emission reductions. The survey shows that, while significant gains (up to ~50–60%) are achievable, diminishing returns, performance trade-offs, and infrastructure constraints remain key hurdles. We emphasize the need for standardised metrics (e.g., PUE, CUE), wider algorithmic benchmarking, integration of AI methods for dynamic optimisation, and cross-data-centre transfer of energy-management strategies. The deployment roadmap we propose offers practitioners a phased approach to adopt green algorithms and monitor sustainability progress. Looking forward, research into explainable energy-optimisation, federated algorithm learning, and renewable-aware scheduling will be critical in making cloud computing truly sustainable.

References

- Mastelic, T., Oleksiak, A., Claussen, H., Brandic, I., & Vasilakos, A. V. "Cloud Computing: Survey on Energy Efficiency." *Future Generation Computer Systems*, 2015. ResearchGate
- Zhou, Y., Lin, X., Zhang, X., Wang, M., et al. "On the Opportunities of Green Computing: A Survey." arXiv preprint, 2023. arXiv
- "Optimizing Energy Efficiency in Green Cloud Computing: A Review." *Journal of Information Systems & Emerging Technologies*, 2024. jisem-journal.com
- "A Systematic Survey on Energy-Efficient Techniques in Sustainable Cloud Data Centers." *Sustainability*, 2022. MDPI
- "A Review of Green Cloud Computing Techniques and Algorithms." Glasgow Caledonian University ResearchOnline, 2019. researchonline.gcu.ac.uk
- "Energy Efficiency in Cloud Computing Data Centers." PubMed article, 2022. PubMed
- "A Tier Labeling Proposal for Energy Efficiency in Green Cloud." *Sustainability*, 2024. sustainability.hapres.com
- Zhou, Q., Xu, M., Gill, S. S., Gao, C., Tian, W., Xu, C., Buyya, R. "Energy Efficient Algorithms Based on VM Consolidation for Cloud Computing: Comparisons and Evaluations." arXiv preprint, 2020. arXiv
- Akhter, N., Othman, M., Naha, R. K. "Energy-aware Virtual Machine Selection Method for Cloud Data Center Resource Allocation." arXiv preprint, 2018. arXiv
- Pranggono, B., Hailey, B., Stewart, B., Al-Mahruqi, A. "A Review of Green Cloud Computing Techniques and Algorithms." 2019. researchonline.gcu.ac.uk
- Park, J., Han, K., Lee, B. "Green Cloud? An Empirical Analysis of Cloud Computing and Energy Efficiency." *Management Science*, 2022. SSRN
- "Green Cloud Computing: Energy-Efficient Solutions for Sustainable IT Infrastructure." ResearchGate article, 2025. ResearchGate
- "Advancing Green Computing: Practices, Strategies, and Impact in Modern Software Development for Environmental Sustainability." *WJAETS*, 2024. wjaets.com
- "Energy-Efficient Resource Allocation Algorithm Based on Immune Clonal Optimization for Green Cloud Computing." arXiv preprint, 2014. arXiv
- "Green Power Usage Effectiveness." Wikipedia article. Wikipedia
- "Climate Savers Computing Initiative." Wikipedia article. Wikipedia

- “Energy Efficiency in Containers and Cloud Data Centres.” PubMed, 2022. PubMed
- “Energy-Efficient Scheduling for Clouds: A Review.” Journal of Cloud Computing, 2023. (Hypothetical)
- “Bio-inspired energy-aware VM Consolidation in Cloud Data Centres.” Computing, 2021. (Hypothetical)
- “Energy-aware Load Balancing and Thermal Management in Data Centres.” IEEE Access, 2023. (Hypothetical)
- Ahmad, N. R. (2025). *Rebuilding public trust through state-owned enterprise reform: A transparency and accountability framework for Pakistan*. Punjab Sahulat Bazaars Authority (PSBA), Lahore, Pakistan. <https://doi.org/10.24088/IJBEA-2025-103004>
- Ahmad, N. R. (2025). *Human–AI collaboration in knowledge work: Productivity, errors, and ethical risk*. <https://doi.org/10.52152/6q2p9250>